

Claim 1 (Currently Amended) A method for constructing and caching a chain of file identifiers that represent a full path to a file system resource comprising the steps of:

processing a file system resource's defined name (DN) into a file identifier (FID) and defined name database;

retrieving the a file identifier for the file system resource that corresponds to the processed defined name of the file system resource, this file identifier being the target file identifier in the chain;

retrieving the file identifier for the next file system resource, said next file resource being the parent of the previous file system resource in the full path;

adding the retrieved file identifier to the chain; and

repeating said retrieving the file identifier for the next file system resource step and said adding the retrieved file identifier to the chain step until a file identifier for each system resource in the full path of the initial file system resource in the chain.

Claim 2 (Previously Amended) The method as described in claim 1 further comprising after said repeating step the steps of:

retrieving a file identifier corresponding to the file system resource, which is the target of the access attempt and a chain file identifier representing the full path directory of the target system resource;

searching for the effective security classification category and defined name for the target resource file identifier;

updating the security classification system, when said search finds a security classification category for the target resource file identifier;

determining whether operations for the target file system resource could affect the file system name space; and

terminating said method when operation does not affect the file system name space.

Claim 3 (Previously Amended) The method as described in claim 2 further comprising the step of flushing the a file identifier chain cache when there is a determination that desired operations on the target file system resource could affect the file system name space.

Claim 4 (Previously Amended) The method as described in claim 2 further comprising before said file identifier (FID) retrieval step the step of processing a system resources defined name (DN) and security classification category into a mapping database which holds a FID to DN mapping.

Claim 5 (Original) The method as described in claim 4 wherein said database-processing step comprises:

providing the defined name and security classification category as inputs;
obtaining a file identifier (FID) for the defined name; and
adding the FID to DN mapping containing the security classification category to the mapping database.

Claim 6 (Previously Amended) The method as described in claim 2 wherein said searching step comprises:

searching the FID to DN mapping database for the security classification category for the FID of the target resource; and

returning the security classification category and defined name for the target FID, when a security classification category for the target FID was found during said search.

Claim 7 (Currently Amended) The method as described in claim 2 wherein said searching step comprises:

- searching the FID to DN mapping database for the security classification category for the FID of the target resource;

- retrieving a FID from the FID chain, when the search does not find a security classification category for the FID of the target resource;

- searching the FID to DN mapping database for the security classification category for the FID of the FID chain; and

- returning the security classification category and defined name for the target FID, when a security classification category for the target FID was found during said search.

Claim 8 (Original) The method as described in claim 7 further comprising the steps of:

- determining whether more entries in the FID chain, when the search does not find a security classification category for the FID used in the search;

- retrieving the next FID in the FID chain; and

- searching the FID to DN mapping database for the security classification category for the currently retrieved FID of the FID chain.

Claim 9 (Original) The method as described in claim 8 further comprising the step of terminating the method when no security classification category is found for any FID in the FID chain.

Claim 10 (Original) The method as described in claim 3 wherein said flushing step comprises:

- retrieving the path name for the target resource, said path name being to a directory for the target resource;

- obtaining a vnode for the directory;

- generating a FID for the directory using the vnode;

- searching for FID chain matching directory FID; and

- removing FID chain from cache, when matching FID chain is found.

Claim 11 (Original) The method as described in claim 10 further comprising before said searching step the step of sorting the FID chains in the FID chain cache into hash list.

Claim 12 (Original) The method as described in claim 11 wherein said searching step comprises: retrieving the first FID chain in the FID chain list;

comparing each FID in said first FID chain to said directory FID;

determining whether there are more FID chains in the list, when said FID chain did not match said directory FID;

retrieving the next FID chain in the FID, and

returning to said comparing step using newly retrieved FID chain.

Claim 13 (Original) The method as described in claim 11 wherein said searching step comprises: retrieving the first FID chain in the FID chain list;

comparing each FID in said first FID chain to said directory FID;

determining whether there are more FID chains in the list, when said FID chain did not match said directory FID; and

terminating method when no FID chain is found.

Claim 14 (Currently Amended) A computer program product in a computer readable medium for use in constructing and caching a chain of file identifiers that represent a full path to a file system resource comprising:

instructions for processing a file system resource's defined name (DN) into a file identifier (FID) and defined name database;

instructions for retrieving the a file identifier for the file system resource that corresponds to the processed defined name of the file system resource, this file identifier being the target file identifier in the chain;

instructions for retrieving the file identifier for the next file system resource, said next file resource being the parent of the previous file system resource in the full path;

instructions for adding the retrieved file identifier to the chain; and

instructions for repeating said retrieving the file identifier for the next file system resource step and said adding the retrieved file identifier to the chain step until a file identifier for each system resource in the full path of the initial file system resource in the chain.

Claim 15 (Previously Amended) The computer program product as described in claim 14 further comprising instructions for retrieving a file identifier corresponding to the file system resource, which is the target of the access attempt and a file identifier chain for the directory of the target system resource;

searching for the effective security classification category and defined name for the target resource file identifier;

updating the security classification system, when said search finds a security classification category for the target resource file identifier;

determining whether operations for the target file system resource could affect the file system name space;

terminating said method when operation does not affect the file system name space; and

flushing the a file identifier chain cache when there is a determination that desired operations on the target file system resource could affect the file system name space.

Claim 16 (Previously Amended) The computer program product as described in claim 15 wherein said flushing instructions comprise:

instructions for retrieving the path name for the target resource, said path name being to a directory for the target resource;

instructions for obtaining a vnode for the directory;

instructions for generating a FID for the directory using the vnode;

instructions for searching for FID chain matching directory FID; and

instructions for removing FID chain from cache, when matching FID chain is found.

Claim 17 (Previously Amended) The computer program product as described in claim 15 wherein said searching instruction further comprises:

instructions for searching the FID to DN mapping database for the security classification category for the FID of the target resource;

instructions for retrieving a FID from the FID chain, when the search does not find a security classification category for the FID of the target resource;

instructions for searching the FID to DN mapping database for the security classification category for the FID of the FID chain; and

instructions for returning the security classification category and defined name for the target FID, when a security classification category for the target FID was found during said search.

Claim 18 (Previously Amended) The computer program product as described in claim 17 further comprising the steps of:

instructions for determining whether more entries in the FID chain, when the search does not find a security classification category for the FID used in the search;

instructions for retrieving the next FID in the FID chain; and

instructions for searching the FID to DN mapping database for the security classification category for the currently retrieved FID of the FID chain.

Claim 19 (Original) The computer program product as described in claim 18 further comprising before said searching, instructions for sorting the FID chains in the FID chain cache into hash list.

Claim 20 (Original) The computer program product as described in claim 19 wherein said searching instructions comprise:

instructions for retrieving the path name for the target resource, said path name being to a directory for the target resource;

instructions for obtaining a vnode for the directory;

instructions for generating a FID for the directory using the vnode;

instructions for searching for FID chain matching directory FID; and

instructions for removing FID chain from cache, when matching FID chain is found.

Claim 21 (Previously Amended) The method as described in claim 2 wherein said file identifier retrieval step comprises:

retrieving the path name of the file resource, which is the target of the access attempt;

obtaining a FID for target resource with said path name;

determining whether obtained FID is in a FID chain; and

returning the target FID and FID chain, when the target resource FID was found in the FID Chain Cache.

Claim 22 (Original) The method as described in claim in further comprising after said path name retrieval step, the step of obtaining vnodes for the target path and parent directory.

Claim 23 (Previously Amended) The method as described in claim 2 wherein said file identifier retrieval step comprises:

retrieving the path name of the file resource, which is the target of the access attempt;

obtaining a FID for target resource with said path name;

determining whether obtained FID is in a FID chain; and

constructing a FID chain for the parent directory, when no FID chain is found.

Claim 24 (Previously Amended) The method as described in claim 23 wherein said FID chain construction comprises:

- setting a temporary vnode to equal the vnode for the parent of the target resource;
- determining whether the temporary vnode is the root directory; and
- inserting FID chain into FID chain into FID chain cache with the first FID in the chain serving as the entry search key, when temporary vnode is the root directory.

Claim 25 (Previously Amended) The method as described in claim 23 wherein said FID chain construction comprises:

- setting a temporary vnode to equal the vnode for the parent of the target resource;
- determining whether the temporary vnode is the root directory;
- retrieving a vnode for the next parent in the directory path and determining whether that parent is the root directory; and
- repeating said retrieving step until parent is the root of the directory.

Claim 26 (Original) The method as described in claim 25 further comprising the step of inserting a completed FID chain into the FID chain cache when the parent is the root directory.

Claim 27 (Canceled)